

Gradient optimization of analytic controls: the route to high accuracy quantum optimal control

Shai Machnes* and David J. Tannor
Weizmann Institute of Science, 76100 Rehovot, Israel

Frank K. Wilhelm and Elie Assémat
Theoretical Physics, Saarland University, Campus, 66123 Saarbrücken, Germany

We argue that quantum optimal control can and should be done with analytic control functions, in the vast majority of applications. First, we show that discretizing continuous control functions as piecewise-constant functions prevents high accuracy optimization at reasonable computational costs. Second, we argue that the number of control parameters required is on-par with the dimension of the object manipulated, and therefore one may choose parametrization by other considerations, e.g. experimental suitability and the potential for physical insight into the optimized pulse. Third, we note that optimal control algorithms which make use of the gradient of the goal function with respect to control parameters are generally faster and reach higher final accuracies than non gradient-based methods. Thus, if the gradient can be efficiently computed, it should be used. Fourth, we present a novel way of computing the gradient based on an equation of motion for the gradient, which we evolve in time by the Taylor expansion of the propagator. This allows one to calculate any physically relevant analytic controls to arbitrarily high precision. The combination of the above techniques is GOAT (Gradient Optimization of Analytic conTrols) - gradient-based optimal control for analytic control functions, utilizing exact evolution in time of the derivative of the propagator with respect to arbitrary control parameters.

I. INTRODUCTION

The ability to drive a quantum system to a desired target in a fast and efficient manner is at the heart of emerging quantum technologies. The task of finding the optimal control pulse to transfer the quantum system to a desired state or to generate a desired quantum gate, has been the subject of extensive research since the first applications of optimal control algorithms [1–3]. Quantum optimal control [4] was originally applied to chemical applications [5], where progress has continued to this day [6]. It has also been applied to nuclear magnetic resonance [7, 8], with applications to medical imaging and spectroscopy. Over the years, the birth of new experimental methods to control quantum systems has led to an increase in the interest in optimal control, which has been applied to processes as diverse as high harmonics generation [9], control of energy flow in biomolecules [10], attosecond physics [11] and quantum computing [12, 13].

Recently, the precision of experimental quantum control has drastically improved, driven primarily by the requirements for quantum computing. Quantum circuits based on Josephson junctions can now approach 10^{-4} error [14] and ion trap architecture can reach comparable precision [15]. Since an even higher accuracy may be necessary for fault tolerant quantum computing [16, 17], one can expect further refinement of the experimental techniques in the near future.

However, no current quantum control algorithm can efficiently achieve very high accuracy, i.e. error below 10^{-6} . The barriers are two-fold. First is the use of a piece-

wise constant (PWC) approximation to smooth control pulses. Quantum control algorithms such as Krotov [2] and GRAPE [7] rely on this approximation, which either compromises accuracy or requires egregious computing resources to achieve high accuracy, as will be shown in the following section. The second barrier is the inherent difficulty of non-gradient based methods to obtain high accuracy. The term non-gradient refers here to genetic algorithms and simplex methods, e.g. the Nelder-Mead algorithm used in the CRAB approach [18]. Thus we argue that if the gradient information can be obtained efficiently, it should be used.

Moreover, quantum optimal control is, at its core, a powerful framework for designing pulses for use in experiments. Hence, it should endeavor to provide control pulses which have a straightforward physical interpretation, and are easy to implement experimentally. The common approach is to define an ansatz characterized by a few parameters which can be tuned in the experiments [19]. The DRAG method [20], for example, was developed for this purpose and is currently used in state of the art experiments on superconducting qubit systems [21]. However, current optimal control algorithms are not designed to work efficiently when the pulse has an analytic parametrization.

To answer these challenges, we present a new approach, Gradient Optimization of Analytical conTrols (GOAT), that avoids the use of a piecewise constant approximation. It is based on formulating an equation of motion for the gradient of the propagator with respect to the control parameters, which is coupled to the Schrödinger equa-

tion. Moreover, we present an efficient method to evolve in time the gradient using a Taylor series expansion of the propagator [22]. GOAT then uses the gradient information to achieve very high accuracy, for any smooth ansatz. After introducing GOAT, we present applications to different two qubit systems, and finally discuss several additional advantages of this approach.

II. OPTIMAL CONTROL PROBLEM

The task of locating the optimal control fields may be viewed as the search for a global optimum of a goal function within a high-dimensional control landscape. Given a drift Hamiltonian H_0 and a number of control Hamiltonians H_k , we endeavor to find the time-dependent *control functions* $c(t)$, parameterized by the parameters $\bar{\alpha}$, such that the overall Hamiltonian,

$$H(\bar{\alpha}, t) = H_0 + \sum_{k=1}^C c_k(\bar{\alpha}, t) H_k \quad (1)$$

drives the system to the desired outcome (e.g. implements a CNOT gate). The control functions are often parameterized as a combination of analytic basis functions, e.g. a Fourier decomposition

$$c_k(\bar{\alpha}, t) = \sum_{j=1}^m A_{k,m} \sin(\omega_{k,m} t + \varphi_{k,m}), \quad (2)$$

or a piecewise-constant (PWC) function. In the Fourier example the parametrization $\bar{\alpha}$ shall be defined as

$$\bar{\alpha} = \{A_{k,m}, \omega_{k,m}, \varphi_{k,m}\}_{k=1\dots C, j=1\dots m}. \quad (3)$$

In the PWC example, the parametrization shall be the value of the controls at each constant piece. The propagator at the final time, $U(T)$, is defined by the time ordered (\mathbb{T}) evolution of the time-dependent Hamiltonian,

$$U(\bar{\alpha}, T) = \mathbb{T} \exp \left(\int_0^T -\frac{i}{\hbar} H(\bar{\alpha}, t) dt \right) \quad (4)$$

The goal function to be minimized is defined as the projective- $S(U)$ distance (infidelity) between the implemented gate and the desired gate, U_{goal} ,

$$g(\bar{\alpha}) := 1 - \frac{1}{\dim(U)} \left| U_{\text{goal}}^\dagger U(T) \right|. \quad (5)$$




Figure 1. Error stemming from piecewise-constant approximation of smooth controls. The time-dependent control (top plot) and the accuracy of propagation when approximated by a PWC function, at various discretizations. The system is composed of 3 qubits with a randomly chosen drift Hamiltonian, a randomly chosen control Hamiltonian and a randomly chosen goal state (see text). The accuracy of propagation is measured by the relative error in the computed infidelity, as compared to a highly accurate adaptive propagation (Runge-Kutta 4/5 with 10^{-14} relative tolerance). The PWC approximation is implemented by sampling the control field at the beginning of the time-slice (blue and green lines), or sampling at the middle of the time-slice (red and cyan), approximating a piecewise-linear interpolation.

Additional tasks and associated goal functions may be defined, such as state transfer problems, open system dynamics, etc. See [23] for further examples.

III. APPROXIMATING CONTINUOUS CONTROLS USING PIECEWISE-CONSTANT FUNCTIONS SHOULD BE AVOIDED

An optimal control procedure utilizing PWC controls will usually converge to machine accuracy with a even a relatively small number of time-slices. However, an experimental implementation that smooths over the PWC control fields will achieve results far below the anticipated accuracy. This is a result of the discrepancy between the propagations induced by the optimized vs. the implemented controls. To estimate this discrepancy, let us consider a simple continuous control, of the form of eq. 2, as depicted in fig. 1 (top). We measure the accuracy of the PWC approximation by comparing it to a highly accurate reference propagation. The results are depicted in fig. 1 (bottom). As one can see, in order to achieve high accuracy, many thousands of segments are required. Specifically, to achieve 10^{-8} accuracy, 10^5 time-slices are required (using middle-of-slice sampling of the control field). The inadequacy of the PWC approximation to the

task of propagating time-dependent Hamiltonians is well known, and has led to the development of more suitable alternatives [24–26].

Such a large number of time-slices is problematic for two reasons. First, it requires an optimization in a control space of very high dimension, many orders of magnitude above what is required, as discussed in the following section. Optimizing in such huge space is unnecessary and wasteful. Second, the optimized PWC control may include high-frequency components, which may result in significant discrepancies when the PWC control is implemented experimentally as a continuous control. We therefore conclude that in cases where smooth controls are implemented experimentally, and high accuracies are desired, PWC approximation of continuous controls are prone to inefficient and imprecise propagation, and hence optimization.

IV. CONTROL SPACE DIMENSION

We now turn to several interrelated issues concerning the analytic parametrization of the control field: How does one decide on a parametrization and its dimensionality, and how does one minimize the possibility of introducing local traps in optimization process? We propose to guide the choice of parametrization by the following guidelines: First, the parametrization should be chosen to fit the experimental setup and/or to offer insight into the mechanism of the optimized pulse. Second, the dimension of the control space must be larger than the dimension of the manifold supporting the dynamics. And third, that while any parameterizations of finite dimension is a form of constraint on the possible controls, the problem of local traps can be significantly reduced using insight into optimization landscape topology [27].

Let us begin by addressing the second issue: dimension of the control space. As is seen in fig. 2a, once the control dimension reaches the dynamic Hilbert space dimension (blue vertical line), optimization virtually always succeeds (color of dots indicate probability of convergence, low probability (red) to the left of the blue line, virtual certainty (green) to the right). It has been our experience that adding a few additional dimensions can help in cases where local traps are abundant (e.g. near the quantum speed limit).

A more in-depth analysis of this issue is visible in fig. 2b, where we clearly see that for state transfer tasks over a Hilbert space of dimension d , a control space of dimension $2d - 2$ is required, while for gate generation tasks a d^2 dimension is required. The data strongly suggests that the match between Hilbert space dimension and minimal control-space dimension holds true regardless of parametrization: Squares denote a simple PWC

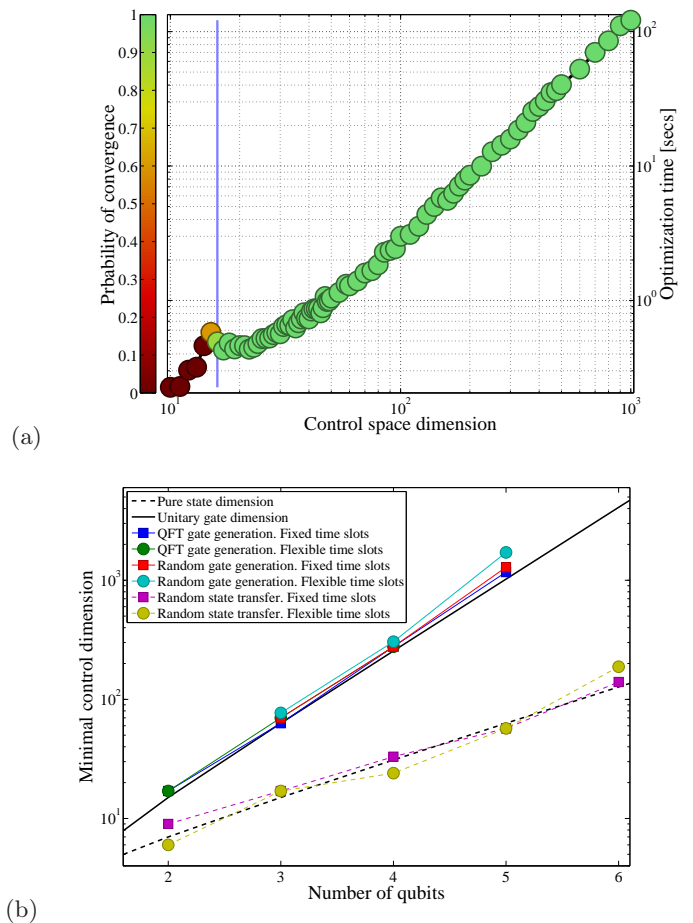


Figure 2. (a) Convergence as a function of control space dimension. Color of dots indicates probability of convergence to fidelity better than 10^{-12} , when starting at a random initial point. Y-axis indicates computation time required for convergence (logarithmic scale). Goal is generation of a 2 qubit random gate (with propagator Hilbert size of 16, blue line), using a single random control Hamiltonian. The control function is parameterized by the Fourier basis, with the control space comprising the amplitudes, and random fixed frequencies and phases. (b) The minimal control space dimension required to always succeed in a random-to-random state transfer task and a random gate generation task, for various system dimensions. Note excellent match of the minimal control space dimension (dots) to the Hilbert space dimension (dashed and solid black lines).

parametrization with a single control dimension per time-slice, while circles denote a PWC parametrization with flexible-width time-slices and two control dimensions per slice. Reference [28] presents data in support of this hypothesis using Fourier parametrization for state-transfer tasks. This makes a strong case that the number of control parameters is independent of the optimal control task details.

Returning to figure 2a, we observe that when one uses significantly higher number of controls than is necessary,

the computational effort required to optimize rises significantly. This is due to both the added numerical effort in calculating gradients as well as the need to search a much larger space, requiring many more iterations and overcoming the many zero-gradient redundant directions. It is therefore clear that one can and should aim for a modest control space dimension, in-line with the size of the dynamic Hilbert space. Note that this strengthens the case against PWC approximation of smooth controls, where the number of control parameters is tied to the number of time slices in the discretization.

Any finite-dimensional parametrization is, by construction, a form of constraint over the infinite-dimensional function space. However, it has been our experience that, with some exceptions, local traps are not a serious problem, and may be mitigated by slightly expanding the control space or alternating between different subspaces of a significantly-expanded control space (e.g. in the Fourier basis we can alternate between optimizing the amplitudes and the frequencies [27]). Additional factors affecting the appearance of local minima include the specifics of the drift and control Hamiltonians, the proximity to the quantum speed limit and the distance of the initial search point relative to the global optimum. The above observations leave us with a great deal of freedom to choose the parametrization of the control fields according to other considerations. Specifically, one would choose a parametrization which fits the experimental constraints (e.g. bandwidth-limited controls which smoothly start and finish at zero). Alternatively, one may choose a parametrization which holds the promise of producing an optimal control sequence from which physical insight can be derived (e.g. Gaussian pulses). The GOAT method, described below, can easily handle any form of constraints and parametrization which may be described by a piecewise-analytic function.

V. GRADIENT OPTIMIZATION OF ANALYTIC CONTROLS (GOAT)

We have seen that smooth analytic controls can have a significant advantage over PWC approximations in accuracy, optimization time, control space dimension and flexibility to match experimental constraints. This stems from the inherent inaccuracy of propagating time-dependent Hamiltonians using PWC approximations, and therefore holds relative to all previous methods that require a PWC approximation. Attempts to enforce analytical control shapes on PWC controls, using the Jacobian matrix to convert PWC gradients to the original analytical parametrization [19, 29], suffer from the same disadvantage - the underlying PWC approximation of smooth controls. What now remains to be shown is that gradient-based optimization of analytic controls is

indeed possible and efficient.

Consider a goal function $g = \text{Tr} \left(U_{\text{goal}}^\dagger U(\bar{\alpha}, t) \right)$. Then the gradient of the goal function with respect to the control parameters is

$$\partial_{\bar{\alpha}} g(\bar{\alpha}) = -\text{real} \left(\frac{g^*}{|g|} \frac{1}{\dim(U)} \text{Tr} \left(U_{\text{goal}}^\dagger \partial_{\bar{\alpha}} U(\bar{\alpha}, t) \right) \right) \quad (6)$$

where $\partial_{\bar{\alpha}} U(\bar{\alpha}, t)$ is the gradient of the propagator with respect to the controls (for details see [23]). The focus then becomes the propagator gradient, $\partial_{\bar{\alpha}} U(\bar{\alpha}, T)$, which can be computed in the following way: The equation of motion for the propagator is

$$\partial_t U(\bar{\alpha}, t) = -\frac{i}{\hbar} H(\bar{\alpha}, t) U(\bar{\alpha}, t). \quad (7)$$

Therefore, the equation of motion for $\partial_{\bar{\alpha}} U$ is simply

$$\partial_t \partial_{\bar{\alpha}} U = -\frac{i}{\hbar} ((\partial_{\bar{\alpha}} H) U + H \partial_{\bar{\alpha}} U), \quad (8)$$

where we have dropped the $(\bar{\alpha}, t)$ dependence for convenience, where $\partial_{\bar{\alpha}} H(\bar{\alpha}, t) = \sum_{k=1}^C (\partial_{\bar{\alpha}} c_k(\bar{\alpha}, t)) H_k$ and where $\partial_{\bar{\alpha}} c_k(\bar{\alpha}, t)$ is parametrization-specific, following eq. 1. See also [30]. Thus, we have two coupled equations of motion, for U and $\partial_{\bar{\alpha}} U$, which may be propagated together as a single aggregate object

$$\partial_t \begin{pmatrix} U \\ \partial_{\bar{\alpha}} U \end{pmatrix} = -\frac{i}{\hbar} \begin{pmatrix} H & 0 \\ \partial_{\bar{\alpha}} H & H \end{pmatrix} \begin{pmatrix} U \\ \partial_{\bar{\alpha}} U \end{pmatrix}. \quad (9)$$

The evolution in time of the propagator gradient may be performed by any mechanism for ODE integration that is accurate and efficient for time-dependent Hamiltonians. Given a parametrization $\vec{\alpha}$ of dimension a , one will need to propagate $a + 1$ objects to compute the full gradient of the goal function. Note that the computational cost of a joint propagation of all components concurrently is significantly cheaper than evaluating each component individually, and moreover may be easily parallelized.

A good choice of propagator, given that we have access to the analytic form of the controls, and hence to all their derivatives, is the Taylor expansion of $U(T)$ at $t = 0$, following [22, 25]. Specifically,

$$U(\bar{\alpha}, T) = \sum_{k=0}^{\infty} \frac{T^k}{k!} \partial_t^k U(\bar{\alpha}, t)|_{t=0} \quad (10)$$

$$\begin{aligned}
\partial_t^0 U &= U(t) \\
\partial_t U &= -\frac{i}{\hbar} H U \\
\partial_t^2 U &= -\frac{i}{\hbar} (\partial_t H U + H(t) \partial_t U) \\
&\vdots \\
\partial_t^k U &= -\frac{i}{\hbar} \sum_{m=0}^{k-1} \binom{k-1}{m} \left(\partial_t^{k-(m+1)} H \right) \partial_t^m U
\end{aligned}$$

Equation 10 can be derived with respect to the control parameters, and as partial differentiation commute, we may write

$$\begin{aligned}
\partial_t^k (\partial_{\bar{\alpha}} U) &= -\frac{i}{\hbar} \sum_{m=0}^{k-1} \binom{k-1}{m} \left(\left(\partial_t^{k-(m+1)} \partial_{\bar{\alpha}} H \right) \partial_t^m U \right. \\
&\quad \left. + \left(\partial_t^{k-(m+1)} H \right) \partial_t^m \partial_{\bar{\alpha}} U \right). \quad (11)
\end{aligned}$$

We have found that evolution in time using this hierarchy of equations is more efficient than Runge-Kutta.

If a single Taylor expansion is not enough to span the entire evolution duration, the result of the previous propagation at $U(t+dt)$ serves as the initial value for the subsequent Taylor expansion. This relay-race mechanism may also be used at control discontinuities in conjuncture with other ODE integration scheme, e.g. Runge-Kutta, allowing GOAT to implement piecewise-analytic parametrization, such as B-splines.

Once one can efficiently compute the goal function derivative with respect to the control parameters, any gradient-driven optimization routine may be used, e.g. Newton methods.

VI. EXAMPLES

Fig. 3 shows the convergence of a CNOT quantum gate generation on a two qubit Ising chain. The system details correspond to Problem 1 of [23]. The ansatz used is a truncated Fourier series with 16 terms. GOAT and Nelder-Mead optimize here exactly the same analytic function, with the same initial guess of the parameters. One can observe that gradient-based methods are indeed preferable to non-gradient methods for high accuracy optimization, when the gradient is available.

Fig. 4 shows a CNOT generation on a two qubit NV center system. Each of the two controls has 5 Fourier components. The system details, including the normalization of time, correspond to Problem 15 of [23]. GOAT




Figure 3. The relative convergence behavior of GOAT (with Newton search) vs. Nelder-Mead. The goal is the generation of a CNOT gate on a two qubit Ising chain. With Nelder-Mead each iteration requires the evaluation of the goal function. With GOAT, each iteration requires the joint evaluation of the goal function and a derivatives, with a being the parametrization dimension. Note that the computational cost is less than $a+1$, as joint propagation is significantly cheaper than evaluating each component individually. No parallelization was used. We therefore elected to measure performance in time required to perform the computation. Note the performance ratio between the two methods should remain constant regardless of hardware.

can reach 10^{-12} accuracy with a smooth control, which would not be feasible with standard algorithms.

VII. DISCUSSION

As with many other gradient-based optimization schemes, GOAT can be used to optimize open quantum systems or goals defined on projected subsystems. The GOAT approach, however, holds several important advantages, in addition to accuracy and adaptability to experimental constraints and intuition with respect to physical mechanisms. These can be divided into the fundamental and the technical.

On the fundamental side, constraints may be implemented by mapping the optimization from an unconstrained space to a constrained subspace, and computing the gradient of the new goal function with the chain rule (e.g. bandwidth and amplitude constraints for a frequency basis or initial and final values for the control fields). Therefore we consider an unconstrained optimization in a constrained subspace [31]. Second, since GOAT provides a convenient way to calculate the derivatives of the goal function with respect to the control parame-




Figure 4. (a) Two control pulses to generate a CNOT gate on NV center system. (b) Timing of the convergence of GOAT toward the optimal solution.

ters, one may extend GOAT to calculating the Hessian. Third, GOAT allows one to produce pulses which are robust to changes in system characterization (e.g. the drift Hamiltonian), by adding the norm of the gradient with respect to the system parameters to the goal function. Fourth, a fundamental property of GOAT is the flexibility of parametrization, allowing one to switch parameterizations mid-way through an optimization. This could be used, for example, to avoid local traps or to accelerate the convergence. Fifth, by specifying equations of motion and not using the Pontryagin's minimum principle, GOAT removes the need for an adjoint state, and its associated backward propagation. This could be very useful in systems where the dynamics is not reversible. Finally, we note the equations of motion describing a device response function (e.g. bandpass filter) may be naturally integrated into the GOAT equations of motion, allowing one to account very accurately for experimental limitations [32].

On the technical side, the propagation of each gradient component does not depend on the propagation of the other gradient components, allowing for parallelization of the computation. Second, all high-order derivatives which are needed for the Taylor expansion of Eq. (11), can be pre-computed automatically using a symbolic algebra engine. Finally, GOAT is memory efficient, as one does not need to store information on the gradient at each time step to compute the goal function. Only the value at the final time is required.

VIII. CONCLUSION

In conclusion, we have presented GOAT - gradient-based optimal control for analytic control functions. GOAT avoids the problematic PWC approximation of time-dependent Hamiltonians, by using a set of novel equations of motion for the gradient of the propagator with respect to the control parameters. We believe GOAT is the method of choice for designing optimal pulses for experiments requiring smooth controls and high accuracy, geared toward high precision quantum technology experiments.

GOAT possesses many additional advantages, mentioned above, which will be explored in depth in the near future. An open-source implementation of GOAT (in Matlab and QuTIP Python toolbox) is currently in development. The authors are interested in pursuing experimental and theoretical applications of the GOAT approach, and will gladly cooperate with other interested parties towards this goal.

* shai.machnes@gmail.com

- [1] S. Conolly, D. Nishimura, and A. Macovski, *IEEE Trans. Med. Imaging* **5**, 106 (1986).
- [2] R. Kosloff, S. Rice, P. Gaspard, S. Tersigni, and D. Tanner, *Chem. Phys.* **139**, 201 (1989).
- [3] A. P. Peirce, Mohammed, A. Dahleh, and H. Rabitz, *Phys. Rev. A* **37**, 4950 (1988).
- [4] C. Brif, R. Chakrabarti, and H. Rabitz, *New. J. Phys.* **12**, 075008 (2010).
- [5] T. Brixner, N. H. Damrauer, P. Niklaus, and G. Gerber, *Nature* **414**, 57 (2001).
- [6] L. Levin, W. Skomorowski, L. Rybak, R. Kosloff, C. P. Koch, and Z. Amitay, *Phys. Rev. Lett.* **114**, 233003 (2015).
- [7] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, *J. Magn. Reson.* **172**, 296 (2005).
- [8] M. Lapert, Y. Zhang, M. Braun, S. J. Glaser, and D. Sugny, *Phys. Rev. Lett.* **104**, 083001 (2010).
- [9] R. Bartels, S. Backus, E. Zeek, L. Misoguti, G. Vdovin, I. P. Christov, M. M. Murnane, and H. C. Kapteyn, *Nature* **406**, 164 (2000).
- [10] T. Buckup, T. Lebold, A. Weigel, W. Wohlleben, and M. Motzkus, *J. Photochem. Photobiol. A* **180**, 314 (2006).
- [11] S. Hassler, T. Balciunas, G. Fan., G. Andriukaitis, A. Pugzlys, A. Baltuska, T. Witting, R. Squibb, A. Zaïr, J. W. G. Tisch, J. P. Marangos, and L. E. Chipperfield, *Phys. Rev. X* **4**, 021028 (2014).
- [12] E. Zahedinejad, J. Ghosh, and B. C. Sanders, *Phys. Rev. Lett.* **114**, 200502 (2015).
- [13] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, J. Biamonte, and J. Wrachtrup, *Nat. Comm.* **5**, 3371 (2014).
- [14] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang,

- R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, *Phys. Rev. Lett.* **113**, 220502 (2014).
- [15] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, *Phys. Rev. Lett.* **113**, 220501 (2014).
- [16] G. Kurizki, P. Bertet, Y. Kubo, K. Molmer, D. Petrosyan, P. Rabl, and J. Schmiedmayer, *PNAS* **112**, 3866 (2015).
- [17] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, *ArXiv* (submitted to PRL, to check) , *arXiv:1501.04932v2* (2015).
- [18] P. Doria, T. Calarco, and S. Montangero, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [19] T. E. Skinner and N. I. Gershenzon, *J. Magn. Reson.* **204**, 248 (2010).
- [20] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [21] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo, *Nat. Comm.* **6**, 6983 (2015).
- [22] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, *Phys. Rev. Lett.* **114**, 090502 (2015).
- [23] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, *Phys. Rev. A* **84**, 022305 (2011).
- [24] C. J. Williams, J. Qian, and D. J. Tannor, *The Journal of Chemical Physics* **95**, 1721 (1991).
- [25] D. Lauvergnat, S. Blasco, X. Chapuisat, and A. Nauts, *J. Chem. Phys.* **126**, 204103 (2007).
- [26] H. Tal-Ezer, R. Kosloff, and I. Schaefer, *Journal of Scientific Computing* **53**, 211 (2012).
- [27] N. Rach, M. M. Müller, T. Calarco, and S. Montangero, *ArXiv* , 1506.04601v1 (2015).
- [28] T. Caneva, A. Silva, R. Fazio, S. Lloyd, T. Calarco, and S. Montangero, *Phys. Rev. A* **89**, 042322 (2014).
- [29] S. Meister, J. T. Stockburger, R. Schmidt, and J. Ankerhold, *J. Phys. A: Math. Theor.* **47**, 495002 (2014).
- [30] I. Kuprov and C. T. Rodgers, *J. Chem. Phys.* **131**, 234108 (2009).
- [31] S. Machnes, D. J. Tannor, F. K. Wilhelm, and E. Assémat, to be published (2015).
- [32] I. N. Hincks, C. Granade, T. W. Borneman, and D. G. Cory, *ArXiv* , 1409.8178v1 (2014).